MEASURING UTILIZATION OF INDIVIDUAL COMPONENTS OF CHANNELS

Technical Field

This invention relates, in general, to channel subsystems and, in particular, to measuring utilization of individual components of channels of a channel subsystem.

Background Art

Historically, the determination of channel utilization has been important for capacity planning, problem analysis, as well as for performance monitoring. Previously, the utilization of a channel path has been measured by accumulating the amount of time that a channel is busy during predefined time intervals. These measurements indicate how busy a channel was during each interval, and based on the measurements, work management and planning decisions were made.

With the advent of more complex channels (e.g., FICON (Fibre Connections) channels offered by International Business Machines Corporation), the previous techniques for measuring channel utilization are inadequate. This is because the new channels are considered a collection of resources and an indication of the channel being busy does not indicate in what capacity it is busy. Thus, a need exists for an improved measurement of channel utilization. In particular, a need exists for a capability that indicates what portion of the channel is busy and the particular utilization of that channel portion.

Summary of the Invention

The shortcomings of the prior art are overcome and additional advantages are provided through the provision of a method of determining utilization of channel components of a computing environment. The method includes, for instance, obtaining measurement data for a selected component of a channel, in which the channel includes a plurality of components; and using the measurement data to determine utilization of the selected component.

In one example, the method further includes obtaining one or more operational characteristics of the selected component, and in yet a further example, the one or more operational characteristics are used in determining the utilization of the selected component.

In a further aspect of the present invention, a method of obtaining information associated with channel components of a computing environment is provided. The method includes, for instance, selecting a channel within the computing environment to be monitored, the channel including a plurality of components; and obtaining data on one or more components of the plurality of components.

In one example, the obtaining data includes obtaining one or more operational characteristics of the one or more components. In a further example, the obtaining data includes obtaining measurement data usable in determining utilization of the one or more components.

In yet a further aspect of the present invention, a method of determining utilization of channels of a computing environment, in which the computing environment includes a plurality of logical partitions, is provided. The method includes, for instance, obtaining measurement data for a channel, and using the measurement data to determine utilization of the channel. The measurement data is representative of use of the channel by the logical partition involved in determining the utilization and of use by one or more other logical partitions of the plurality of logical partitions.

System and computer program products corresponding to the above-summarized methods are also described and claimed herein.

Additional features and advantages are realized through the techniques of the present invention. Other embodiments and aspects of the invention are described in detail herein and are considered a part of the claimed invention.

Brief Description of the Drawings

The subject matter which is regarded as the invention is particularly pointed out and distinctly claimed in the claims at the conclusion of the specification. The foregoing and other objects, features, and advantages of the invention are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

FIG. 1 depicts one example of a computing environment incorporating and using the capabilities of the present invention;

FIG. 2a depicts one example of a channel utilization block used in accordance with an aspect of the present invention;

FIG. 2b depicts one example of a channel utilization entry of the channel utilization block of FIG. 2a, in accordance with an aspect of the present invention;

FIG. 3 depicts one example of a channel subsystem extended CPMF data area used in accordance with an aspect of the present invention;

FIG. 4 depicts one embodiment of various components of a channel of FIG. 1, in accordance with an aspect of the present invention;

FIG. 5 depicts one example of a channel processor universal CPMF data area used in accordance with an aspect of the present invention;

FIG. 6 depicts one embodiment of the logic used to set the appropriate mode for a channel processing measurement facility used in accordance with an aspect of the present invention;

FIG. 7 depicts one embodiment of the logic associated with an overview of how to determine

utilization of various components of a channel, in accordance with an aspect of the present invention;

FIG. 8 depicts one embodiment of the logic associated with an operating system determining the utilization of one or more channel components, in accordance with an aspect of the present invention;

FIG. 9a depicts one embodiment of a channel measurement characteristics block used in accordance with an aspect of the present invention;

FIG. 9b depicts one example of information included in the CMG-dependent channel measurement characteristics field of the block of FIG. 9a, in accordance with an aspect of the present invention;

FIG. 10 depicts one example of the command request block for a Set Extended Channel Measurements command used in accordance with an aspect of the present invention;

FIG. 11 depicts one embodiment of the logic associated with a channel processor collecting measurement data for one or more channel components, in accordance with an aspect of the present invention; and

FIG. 12 depicts one embodiment of the logic associated with an I/O processor storing measurement data for one or more channel components in channel utilization blocks, in accordance with an aspect of the present invention.

Best Mode for Carrying Out the Invention

In accordance with an aspect of the present invention, channel utilization is measured. In particular, various individual components of a channel are monitored and measured to determine the utilization of each particular component. Further, in accordance with another aspect of the present invention, when the system is in logical partition mode, the collected measurement data reflects the utilization of the channel (or channel component) by a particular logical partition, as well as the utilization by other logical partitions in the system, which are coupled to the channel (or channel component).

One embodiment of a computing environment incorporating and using the channel measurement capabilities of the present invention is described with reference to FIG. 1. A computing environment 100 is based, for instance, on the Enterprise Systems Architecture (ESA)/390 offered by International Business Machines Corporation, Armonk, New York. ESA/390 is described in an IBM publication entitled "Enterprise Systems Architecture/390 Principles of Operation," IBM Publication No. SA22-7201-06, July 1999, which is hereby incorporated herein by reference in its entirety. One example of a computing environment based on ESA/390 is the 9672 Parallel Enterprise Server offered by International Business Machines Corporation.

Computing environment 100 includes, for example, at least one central processing complex (CPC) 102 having one or more logical partitions (LPARs) 104, a hardware systems area (HSA) 106 and at least one channel subsystem 108.

The resources of the CPC are allocated to the logical partitions. In particular, each logical partition 104 has one or more logical processors (not separately shown for clarity), each of which represents all or a share of a physical processor allocated to the partition. The logical processors of a particular partition may be either dedicated to the partition (so that the underlying processor resource is reserved for that partition) or shared with another partition (so that the underlying processor resources are potentially available to another partition).

In the particular example shown, each of logical partitions 1-M functions as a separate system having a resident operating system 109 (which may differ for each logical partition). In one embodiment, the operating system is the OS/390 or MVS/ESA operating system offered by International Business Machines Corporation.

Further, each partition has a portion of the CPC's main storage resources assigned to it and included within the storage are channel utilization blocks 110. One example of a channel utilization block 110 is described with reference to FIGs. 2a-2b.

Referring to FIG. 2a, channel utilization block 110 includes one or more channel utilization entries 200, each of which contains the utilization information for one channel. In one example, channel utilization entry 200 includes the following fields:

(A) A channel-utilization-information validity (CUIV) map 202 that applies to the contents of the channel utilization entry. The bits

of the CUIV map indicate which words of the channel utilization entry contain valid channel utilization information;

(B) A timestamp 204 that identifies when channel utilization information was last stored in the channel utilization entry; and

(C) CMG-dependent channel-utilization information 206, the contents of which are dependent upon the channel measurement group (CMG) to which the channel belongs. The program can determine the CMG for a channel by executing a Store Channel Measurement Characteristics command (described below) and examining the CMG field in the channel measurement characteristics block (CMCB) (also described below) for the associated channel.

Various examples of values for CMG and the corresponding contents of fields of entry 200 include the following:

CMG=0: A value of 0 in the CMG field indicates that the channel-measurement group for the specified channel path identifier (CHPID) is unknown. The CUIV bits are zeros and the contents of the remainder of the channel-utilization entry have no meaning.

CMG=1: The CMG-dependent channel-utilization information for CMG=1 includes the following, in one example:

(A) Channel-Path-Busy Time (CPBT) (CPC): This field contains the accumulation of all measured channel-path-busy intervals for the associated channel path. The channel path is busy when it is actively communicating with an attached control unit or device.

When the CPC is operating in LPAR mode, the CPBT (CPC) value includes the channel-path-busy intervals associated with all of the logical partitions that have access to the associated channel path.

(B) Channel-Path-Busy Time: When the CPC is operating in LPAR mode, the channel path busy time contains the accumulation of measured channel-path-busy intervals for the associated channel path. The accumulation only includes intervals that are associated with the logical partition that has access to this channel utilization block. The channel path is busy when it is actively communicating with an attached control unit or device.

When the CPC is operating in basic mode (non-LPAR), the contents of this field have no meaning and the appropriate bit of the CUIV field is zero.

CMG-2: The CMG-dependent channel-utilization information for CMG=2 includes the following, in one example:

(A) Count of Bus Cycles (CPC): This field contains an accumulated count of the internal bus cycles that have been used for I/O operations on the associated channel path.

When the CPC is operating in LPAR mode, the count includes the internal bus cycles that were used for I/O operations for all logical partitions that have access to the associated channel path.

(B) Count of Channel Work Units (CPC): This field contains an accumulated count of the number of channel work units that have been performed by the associated channel path.

When the CPC is operating in LPAR mode, the count includes the channel work units that were performed for all logical partitions that have access to the associated channel path.

(C) Count of Channel Work Units: When the CPC is operating in LPAR mode, this field contains an accumulated count of the number of channel work units that have been performed by the associated channel path. The count includes only the channel work units that were performed for the logical partition that has access to this channel utilization block.

When the CPC is operating in basic mode, the contents of this field have no meaning and

the appropriate bit of the CUIV field is zero.

(D) Count of Data Units Written (CPC): This field contains an accumulated count of the number of data units that have been written using the associated channel path.

The channel measurement characteristics block for the associated channel path specifies the number of bytes contained in a data unit.

When the CPC is operating in LPAR mode, the count includes the data units that were written for all logical partitions that have access to the associated channel path.

(E) Count of Data Units Written: When the CPC is operating in LPAR mode, this field contains an accumulated count of the number of data units that have been written using the associated channel path. The count includes only the data units that were written for the logical partition that has access to this channel-utilization block.

When the CPC is operating in basic mode, the contents of this field have no meaning and the appropriate bit of the CUIV field is zero.

(F) Count of Data Units Read (CPC): This field contains an accumulated count of the number

of data units that have been read using the associated channel path. The channel measurement characteristics block for the associated channel path specifies the number of bytes contained in a data unit.

When the CPC is operating in LPAR mode, the count includes the data units that were read for all logical partitions that have access to the associated channel path.

(G) Count of Data Units Read: When the CPC is operating in LPAR mode, this field contains an accumulated count of the number of data units that have been read using the associated channel path. The count includes only the data units that were read for the logical partition that has access to this channel-utilization block.

When the CPC is operating in basic mode, the contents of this field have no meaning and the appropriate bit of the CUIV field is zero.

Returning to FIG. 1, coupled to the logical partitions is a hardware systems area (HSA) 106 which, in this embodiment, includes a channel subsystem extended CPMF data area 112. The channel subsystem extended CPMF data area includes measurement information for individual channels, as described with reference to FIG. 3. In particular, for each channel, the data area includes, for instance, channel measurement characteristics 300, total measurement data 302,

and measurement data 304 for each particular operating system associated with the specified channel. Further details regarding these fields are provided below.

Channel measurement characteristics 300 of a channel include, for example, maximum bus cycles specifying the maximum internal bus cycles per second, available for I/O operations on the specified channel path; maximum channel work units specifying the maximum number of channel work units that can be performed per second by the specified channel paths; maximum write data units specifying the maximum number of data units that can be written per second using the specified channel path; maximum read data units specifying the maximum number of data units that can be read per second using the specified channel path; and data unit size specifying the number of bytes of data that are contained in the data unit for the specified channel path.

Total measurement data 302 includes the total measurements for specified characteristics of a channel. In one example, the total measurement data includes the total bus cycles, the total work units, the total written data units and the total read data units, as explained further below.

Additionally, the channel measurement data per operating system includes the measurements of the specified characteristics on a per operating system basis.

Referring again to FIG. 1, also coupled to the HSA is channel subsystem 108, which is used in communicating with input/output (I/O) devices. Channel subsystem 108 directs the flow of information between the input/output devices and

main storage. It relieves the central processing units of the task of communicating directly with the input/output devices and permits data processing to proceed concurrently with input/output processing.

In one embodiment, channel subsystem 108 includes one or more I/O processors 114 coupled to one or more channels 116. Further, coupled to each channel is a physical channel link 118 that is used to communicate with the I/O devices. In particular, in one example, the link couples the channel to a control unit, which is further coupled to one or more I/O devices.

Further details regarding a channel are described with reference to FIG. 4. In one embodiment, each channel 116 includes, for instance, a channel processor storage 400 coupled to a host-bus bridge and DMA controller 402 for providing access, via a system bus 404, to the operating system and other components of the CPC. Additionally, coupled to the host-bus bridge and DMA controller is an internal bus 406, which is coupled to a channel processor 408 and an adapter card 410. The adapter card provides access to the I/O devices via an external link 412.

One example of a channel is a FICON channel offered by International Business Machines Corporation. One example of a FICON channel is described in "IBM S/390 FICON Implementation Guide", IBM Publication No. SG24-5169-00 (Nov. 1999), which is hereby incorporated herein by reference in its entirety.

In one embodiment, channel processor storage 400 includes a channel processor universal CPMF data area 414

used in storing various measurement information for the channel. As shown in FIG. 5, channel processor universal CPMF data area 412 includes, for instance, channel measurement characteristics 500, total channel measurement data 502, and measurement data 504 for each operating system coupled to the channel.

Channel measurement characteristics data 500 includes, for instance, maximum units 506 for a channel component. For example, maximum units 506 includes the maximum bus cycles (e.g., for Component 1 -- the internal bus) of the channel, the maximum channel work units (e.g., for Component 2 -- the channel processor), the maximum write data units (e.g., for Component 3 -- the external link), and the maximum read data units (e.g., for Component 4 -- the external link) for the channel.

Total channel measurement data 502 includes the total measurements for the components (e.g., the total bus cycles, total work units, total written data units and the total read data units).

Similarly, the operating system measurement data includes an operating system timestamp 514, as well as accumulated units 516 for each operating system of a component. Again, this is the accumulated information for each component.

Described in detail above is one example of a computing environment incorporating and using aspects of the present invention. The above-described environment is only one example, however. Other computing environments can be used without departing from the spirit of the invention.

Further, for one or more aspects of the present invention, the computing environment need not be in logical partitioning (LPAR) mode. Basic mode (non-LPAR) can also benefit from aspects of the present invention.

~~Channel subsystems are further described in Casper et~~ al., U.S. Patent No. 5,526,484, entitled "Method and System for Pipelining the Processing of Channel Command Words," issued on June 11, 1996, as well as in Galbraith et al., U.S. Patent No. 5,265,250, entitled "Channel Measurement Method and Means," issued on November 23, 1993, each of which is hereby incorporated herein by reference in its ~~entirety.~~ A1

In accordance with an aspect of the present invention, various components of the channel are individually measurable. These components include, for instance, internal bus 406, channel processor 408, and external link 412. By measuring these components individually, specific information about the individual components, as well as about the channel as a whole, can be gathered and employed in work management decisions, billing situations, system performance, as well as other areas.

Each of the various components is measured using a channel path measurement facility (CPMF). The channel path measurement facility may be initiated under various conditions, as described with reference to FIG. 6. As one example, the operating system may run the channel path measurement facility when an initial program load (IPL) is performed, STEP 600. As a further example, the operating system runs the channel path measurement facility when an operator selects an SRM parmlib member using an OS/390 Set

command, STEP 602. As a further example, the operating system may run the channel path measurement facility when an operator configures a channel path (CHPID) online, STEP 604.

Regardless of the triggering event for initiating the channel path measurement facility, a determination is made as to whether the parmlib specifies a channel path measurement facility (CPMF) keyword, INQUIRY 606. In particular, a decision is made as to whether the systems programmer chose to run with a parmlib member that specifically indicated the desired CPMF mode (e.g., compatibility mode, CPMF=COMPAT, in which the individual components are not measured; or extended mode, CPMF=EXTENDED, in which the individual components are measured in accordance with aspects of the present invention).

If the CPMF keyword was specified, then the system operates in the mode selected, STEP 608. However, if the systems programmer did not specifically set the CPMF parameter, then the operating system chooses the CPMF mode based on the set of channel path ids found in the machine. In particular, a determination is made as to whether an extended measurement mode CHPID is found, INQUIRY 610. If any of the online CHPIDs desire extended mode for gathering the performance data, then CPFM is initialized in extended mode, STEP 612. Otherwise, if no CHPIDs desiring extended mode are found, the system runs in the old compatibility mode, STEP 614. This completes initialization of the channel path measurement facility in the desired mode, STEP 616.

In one embodiment, both the compatibility mode and the extended mode can be active at the same time. That is, one facility can be activated in one mode, and another facility can be activated in another mode. This is true for both the same logical partition and across multiple logical partitions.

Assuming that the channel path measurement facility is set to extended mode, the channel path measurement facility collects data on the various components of the channel, in accordance with an aspect of the present invention. An overview of one embodiment of the logic employed to collect the various data is described with reference to FIG. 7.

Initially, a Store Channel Measurement Characteristics CHSC command is used by the operating systems to learn about the measurement characteristics (or attributes) of each channel that is to be measured, STEP 700. That is, the operating systems learn about the characteristics of the various components of the channel to be measured. In one example, the measurement characteristics include the maximum bus cycles, which is an attribute of internal bus 406 (FIG. 4); maximum channel work units, which are attributes of channel processor 408; and maximum written/read data units, which are attributes associated with external link 412. Further, the characteristics include the data unit size.

In addition to using the Store Channel Measurement Characteristics CHSC command to learn about the measurement attributes, a Set Extended Channel Measurements CHSC command (described below) is used by the operating system to start the channel path measurement facility and to provide to the

I/O processor a program data area for the channel utilization blocks, STEP 702.

Thereafter, each channel processor collects measurement data, via, for instance, the channel path measurement facility, for the various channel components contained in that channel and places the data in channel processor universal CPMF data area 414 (FIG. 4) contained in channel processor storage 400. Then, at each proper interval, the data is copied from data area 414 to channel subsystem extended CPMF data area 112 (FIG. 1) for that channel, STEP 704.

Additionally, the I/O processor collects from each of the channel subsystem extended CPMF data areas the data for all of the channels, reformats it and stores it in the channel utilization blocks in the previously provided program data area, STEP 706. The operating system then observes the information that has been placed in the channel utilization blocks and is able to interpret the information using the measurement characteristics that apply to each channel that were previously gathered (in STEP 700), STEP 708. This interpreted information provides the utilization of the measured channel components.

Further details regarding the processing performed by the various units of the central processing complex to measure the utilization of the individual components of a channel are described with reference to FIGs. 8-12. In particular, FIG. 8 describes one embodiment of the processing performed by one or more of the operating systems; FIG. 11 depicts one embodiment of the logic performed by one or more of the channel processors; and FIG.

12 depicts one embodiment of the logic performed by one or more of the I/O processors.

Referring to FIG. 8, initially, the operating system collects the channel measurement characteristics data for each channel associated with that operating system, using the Store Channel Measurement Characteristics CHSC command, STEP 800. As one example, the command request block for the store command includes an op code specifying the Store Channel Measurement Characteristics command, as well as an identifier for each channel path (CHPID) for which information is requested.

In response to this request, a command response block for the Store Channel Measurement Characteristics command is provided. This response block includes, for instance, one or more channel measurement characteristics blocks. In particular, a channel measurement characteristics block may be stored for one or more of the specified CHPIDs, beginning with the CHPID specified by the first CHPID field. One example of a channel measurement characteristics block is described with reference to FIG. 9a.

In one example, a channel measurement characteristics block 900 includes, for instance, one or more flags 902, a channel path identifier (CHPID) 904, a channel measurement characteristics validity (CMCV) indicator 906, a channel measurement group (CMG) 908 and a channel measurement group-dependent measurement characteristics field 910, each of which is described below.

Flags 902 provide indicators that enable the proper interpretation of the information in the channel measurement

characteristics block. As one example, flags 902 include a non-valid (N) indicator which, when zero, indicates that information is provided in the channel measurement characteristics block for the associated channel path. The CMG and CMCV fields indicate which information is provided. When one, the non-valid bit indicates that information is not provided for the associated channel path and the CMCV field is all zeros.

Flags 902 also include a shared channel path (S) indicator, which is meaningful when the N bit is zero and the CPC is operating in LPAR mode. When zero, the S bit indicates that the associated channel path is not shared between the logical partition executing the Store Channel Measurement Characteristics command and other logical partitions. When one, the S bit indicates that the associated channel path may be shared between the logical partition executing the Store Channel Measurement Characteristics command and one or more of the logical partitions.

CHPID 904 contains the channel path id of the channel path with which this channel measurement characteristics block is associated.

Channel measurement characteristics validity (CMCV) indicator, when one, indicates that the channel path measurement facility stores channel measurement characteristics information in the corresponding word of the channel measurement characteristics block. When zero, a CMCV bit indicates that the contents of the corresponding word of the channel measurement characteristics block have no meaning.

Channel measurement group (CMG) 908 specifies the measurement group to which the specified CHPID belongs. This value determines the contents and format of the CMG-dependent channel measurement characteristics area, and the format and contents of the channel utilization entry for the associated channel path, in the channel utilization block, described above.

CMG-dependent channel measurement characteristics 910 includes information which is dependent upon the value specified in the channel measurement group field. For instance, a value of zero in the CMG field indicates that the channel measurement group for the specified CHPID is unknown. Thus, the CMCV bits are zeros and the contents of block 910 have no meaning. Further, when CMG=1, the CMCV bits are zeros, and again, the contents of block 910 have no meaning.

However, when the CMG field contains a value of 2, the CMG-dependent channel measurement characteristics include, for instance, the following: maximum bus cycles 912 (FIG. 9b); maximum channel work units 914; maximum write data units 916; maximum read data units 918; and a data unit size 920, each of which is described above.

Returning to FIG. 8, in addition to collecting the channel measurement characteristics data for each channel, the operating system starts the channel path measurement facility using a Set Extended Channel Measurement CHSC command, STEP 802. The Set Extended Channel Measurement command starts or stops the channel facility based on the operation requested in the request block of the command. If the channel path measurement facility is to be started, one

facility is provided, when the CPC is operating in basic mode. However, when the CPC is operating in LPAR mode, one channel path measurement facility is provided for each logical partition. The phrase "the channel path measurement facility" used herein implies the CPC channel path measurement facility, when the CPC is operating in basic mode, or the channel path measurement facility associated with the logical partition executing the command, when the CPC is operating in LPAR mode. One example of the request block used for the Set Extended Channel Measurement command is described with reference to FIG. 10.

In one example, a command request block 1000 for the Set Extended Channel Measurement command includes an operation code 1002 specifying the operation to be performed. In one example, this code may indicate a start of the channel path measurement facility for the Set Extended Channel Measurements command; a stop of the channel path measurement facility for the Set Extended Channel Measurements command; or a test of the status of the channel path measurement facility for the Set Extended Channel Measurements command.

Additionally, the command request block includes a key 1004 used by the channel path measurement facility to access the channel utilization block. Further, it includes utilization block addresses 1006. When the operation code specifies a start operation, the channel utilization block addresses contain the absolute addresses of, for instance, two 4k-byte areas in main storage which collectively make up the channel utilization block. Each part (i.e., 4k area) of the channel utilization block contains utilization information for a unique set of a predefined number (e.g.,

128) of channel paths. When the operation code specifies a stop or a test operation, the contents of the channel utilization block address is ignored.

In response to executing the Set Extended Channel Measurement command, a response block is provided. This response block, indicates, for instance, the results of executing the Set Extended Channel Measurement command.

The channel path measurement facility provides the program (e.g., operating system) with the ability to accumulate channel path utilization information, and to accumulate the portion of shared channel path utilization attributable to a logical partition, when the CPC is operating in LPAR mode.

Returning to FIG. 8, subsequent to starting the channel path measurement facility, the channel subsystem begins to accumulate channel measurement data into the associated channel utilization block at predefined intervals (e.g., at least once every four seconds), STEP 804. One example of how this is performed is described below with reference to FIGs. 11 and 12.

Thereafter, the operating system determines the utilization of a particular channel component (such as, for instance, the channel processor, the internal bus or the external link), which is referred to herein as Component X, STEP 806.

One embodiment of how the operating system performs this determination is described with reference to STEP 808. For example, a copy of the channel utilization block, called

Sample # 1, is saved. Then, after a predefined delay (e.g., 15 seconds), another copy of the channel utilization block, called Sample # 2, is also saved. Subsequently, the "change in channel Component X units" between the sample periods is calculated by subtracting the count of channel Component X units from Sample # 1 from the count of channel Component X units from Sample # 2. (For example, the count of bus cycles for the internal bus of Sample # 1 is subtracted from the count of bus cycles for the internal bus of Sample # 2. Similar computations are made for the various characteristics of the other components.)

Additionally, the "change in channel time" between the sample periods is calculated by subtracting the channel timestamp for Sample # 1 from the channel timestamp for Sample # 2. Thereafter, the "average change in channel Component X units per second" during this sample period is calculated by dividing the change in channel Component X units by the change in channel time. The utilization of channel Component X, during the sample period, is then calculated by dividing the "average change in channel Component X units per second" by the maximum channel component units per second from the channel measurement characteristics.

Subsequent to determining the utilization of a channel Component X, processing proceeds to STEP 804 in order to continue to collect and analyze data for another of the various channel components.

As mentioned above, in addition to the operating system performing tasks relevant to determining the utilization of

channel components, the channel processor also performs various tasks needed to collect measurement data relating to a measurable channel component. One embodiment of the logic associated with collecting this data is described with reference to FIG. 11.

Initially, the channel processor is reset and initialized, STEP 1100. Thereafter, the channel processor stores its channel measurement characteristics (CMC) into the channel subsystem extended CPMF data area in the HSA, STEP 1102. Subsequently, a determination is made as to whether channel work is pending, INQUIRY 1104. In one example, this determination is made by checking an internal status indicator. If no work is pending, then no data need be collected, thus processing continues with INQUIRY 1106.

At INQUIRY 1106, a determination is made as to whether a predefined time period (e.g., at least two seconds) has passed since the HSA channel measurement data has been updated. If the predefined amount of time has not elapsed, then processing continues with INQUIRY 1104. However, if the predefined time interval has passed, then the channel's copy of the measurement data is stored into the HSA, STEP 1108. In particular, the data stored in the channel processor universal CPMF data area is copied to the channel subsystem extended CPMF data area. Thereafter, processing continues with INQUIRY 1104.

At INQUIRY 1104, if channel work is pending, then the requested I/O operation is performed for the specified operating system, STEP 1110. Additionally, measurement data related to the current I/O operation is collected for each measurable channel component, STEP 1112. The new measurement

data is then added to the accumulated measurement data for the particular operating system and to the accumulated total channel measurement data in the channel processor universal CPMF data, STEP 1114. Additionally, the measurement data timestamp is updated. Thereafter, processing continues with INQUIRY 1106.

Processing related to the I/O processor's collection of data is described with reference to FIG. 12. Initially, a set loop count is set equal to a variable M, where M is equal to the number of operating systems, STEP 1200. Thereafter, a determination is made as to whether the extended channel path monitoring facility is active for this operating system, INQUIRY 1202.

If the facility is not active for this operating system, then the loop count is decremented by a predefined amount (e.g., by one), and a determination is made as to whether the last operating system has been processed, INQUIRY 1204. If the last operating system has been processed, then the logic polls for other IOP work to perform, STEP 1206. Otherwise, processing continues with INQUIRY 1202 for the next operating system.

At INQUIRY 1202, if the extended monitoring facility is active for this operating system, then a channel count is set equal to a variable, N, which reflects the number of channels, STEP 1208. Thereafter, the total measurements data and the operating system measurement data for the channels in which the operating system has access are fetched from the channel subsystem extended CPMF data area of the HSA and stored into the operating system's channel utilization blocks, STEP 1210. Subsequently, the variable,

channel count, is decremented by a predefined amount (e.g., one), and a determination is made as to whether the last channel has been processed, INQUIRY 1212. If the last channel has been processed, then processing continues with INQUIRY 1204. Otherwise, processing continues with STEP 1210.

Described in detail above is a measurement facility that measures utilization of individual channel components. This advantageously provides utilization information in enough detail (e.g., work units, data units written, data units read and bus cycles used) to effectively assess the workload of a channel. Advantageously, measurements are provided that are tailored to the type of channel path, and which are meaningful in determining how, and to what degree, the channel is being utilized. The utilization measurements of the present invention include information concerning both the CPC level and the single logical partition level. In addition, in one aspect of the present invention, an operating system that is executing in a logical partition under LPAR is able to observe channel path utilization for all logical partitions, as well as for its own logical partition.

Although various components are described herein, this invention is not limited to those components. Aspects of this invention are universal and can be used with a wide range of components. Further, the characteristics described herein are only examples. Other characteristics can be used without departing from the spirit of the present invention.

The present invention can be included in an article of manufacture (e.g., one or more computer program products)

having, for instance, computer usable media. The media has embodied therein, for instance, computer readable program code means for providing and facilitating the capabilities of the present invention. The article of manufacture can be included as a part of a computer system or sold separately.

Additionally, at least one program storage device readable by a machine, tangibly embodying at least one program of instructions executable by the machine to perform the capabilities of the present invention can be provided.

The flow diagrams depicted herein are just examples. There may be many variations to these diagrams or the steps (or operations) described therein without departing from the spirit of the invention. For instance, the steps may be performed in a differing order, or steps may be added, deleted or modified. All of these variations are considered a part of the claimed invention.

Although preferred embodiments have been depicted and described in detail herein, it will be apparent to those skilled in the relevant art that various modifications, additions, substitutions and the like can be made without departing from the spirit of the invention and these are therefore considered to be within the scope of the invention as defined in the following claims.